Evolving spatio-spectral feature extraction algorithms for hyperspectral imagery

Steven P. Brumby[*] and Amy E. Galbraith

Space and Remote Sensing Sciences, Los Alamos National Laboratory,
Mail Stop D436, Los Alamos, New Mexico 87545, USA.

ABSTRACT

Hyperspectral imagery data sets present an interesting challenge to feature extraction algorithm developers. Beyond the immediate problem of dealing with the sheer amount of spectral information per pixel in a hyperspectral image, the remote sensing scientist must explore a complex algorithm space in which both spatial and spectral signatures may be required to identify a feature of interest. Rather than carry out this algorithm exploration by hand, we are interested in developing learning systems that can evolve these algorithms.

We describe a genetic programming/supervised classifier software system, called GENIE, which evolves image processing tools for remotely sensed imagery. Our primary application has been land-cover classification from satellite imagery. GENIE was developed to evolve classification algorithms for multispectral imagery, and the extension to hyperspectral imagery presents a chance to test a genetic programming system by greatly increasing the complexity of the data under analysis, as well as a chance to find interesting spatio-spectral algorithms for hyperspectral imagery. We demonstrate our system on publicly available imagery from the new Hyperion imaging spectrometer onboard the NASA Earth Observing-1 (EO-1) satellite.

Keywords: Genetic Programming, Hyperspectral Imagery, Feature Extraction, Image Processing, Remote Sensing.

1. INTRODUCTION

Los Alamos National Laboratory's GENIE software[1-4] is a machine learning software system using techniques from the fields of genetic algorithms (GA)[5-7] and genetic programming (GP)[8] to construct feature extraction algorithms for remotely sensed imagery. Both the structure of the feature extraction algorithm, and the parameters of the individual image processing steps, are learned by the system. The algorithms evolved by GENIE combine spatial and spectral processing, and the system was in fact designed to enable exploration of spatio-spectral image processing. This system has been shown to be effective in detecting a range of complex terrain features, such as golf courses in MODIS Airborne Simulator imagery[9], delineating and classifying wildfire burn scars[10] and vegetation land-cover classes[11] using a number of moderate-resolution (Landsat class) multispectral imagery datasets, and more recently, detecting cratered terrain on Mars[12] using a high-resolution panchromatic dataset from the Mars Global Surveyor/Mars Orbital Camera. We now describe work exploring the use of this system with a challenging type of remotely sensed data: hyperspectral imagery.

GENIE follows the paradigm of genetic programming: a population of candidate image-processing algorithms is randomly generated from a collection of low-level image processing operators, including texture measures, spectral band-math operations (e.g. ratios of bands), and morphological filters. The fitness of each individual is assessed from its performance on training data provided by the human user via a graphical interface. A point-and-click interface is used by a human analyst to label example pixels containing the feature of interest, and also label pixels of the background (non-feature). Our fitness metric in based on measuring the total error rate (false positives and false negatives) on the feature extraction task. After a fitness value has been assigned to each candidate algorithm in the population, the most

[*] Further author information: (Send correspondence to S.P.B.) Email: `brumby@lanl.gov`

fit members of the population reproduce with modification via the evolutionary operations of mutation and crossover. This process of evaluation and reproduction is iterated until the population converges, or some desired level of classification performance is attained, or some user-specified limit on computational effort is reached (e.g. number of candidate algorithms evaluated). The final result is a grey-scale enhancement of the feature of interest, which is then converted into a final boolean classification using a threshold. This final threshold may be adjusted by the human user to take into account the desired emphasis of the value of detection rate (true positives) over false alarms and missed detections. The GENIE software is written in a mixture of Perl, Java, and the IDL image processing language, and runs on standard Linux and Solaris workstations. A typical run of 50 generations with a population of 50 candidate solution running on a single workstation, takes on the order of 1 hour to complete.

We now briefly describe our method of providing training data, our encoding of image-processing algorithms as chromosomes for manipulation by the GA, our method for evaluating the fitness of individuals in the population, and the voting scheme adopted. In the next section we describe the results of our experiments with hyperspectral imagery, and we end with some brief conclusions.

2. *GENIE* FEATURE EXTRACTION TECHNIQUE

2.1 Training Data

The environment for the population consists of one or a number of training scenes. Each training scene contains a raw image, together with a weight plane and a truth plane. The weight plane identifies the pixels to be used in training, and the truth plane locates the features of interest in the training data. Providing sufficient quantities of good training data is crucial to the success of any machine learning technique. In principle, the weight and truth planes may be derived from an actual ground campaign (i.e., collected on the ground at the time the image was taken), may be the result of applying some existing algorithm, and/or may be marked-up by hand using the best judgement of an analyst looking at the data. We have developed a graphical user interface (GUI), called Aladdin, for manual marking up of raw imagery. Using Aladdin, the analyst can view a raw image in a variety of ways, and can mark up training data by painting directly on the image using the mouse. Training data is ternary-valued, with the possible values being “true”, “false”, and “unknown”. True defines areas where the analyst is confident that the feature of interest does exist. False defines areas where the analyst is confident that the feature of interest does not exist. Unknown pixels do not influence the fitness of a candidate algorithm.

2.2 Representation of Image-Processing Algorithms

Traditional genetic programming[8] (GP) uses a variable sized (within limits) tree representation for algorithms. Our representation differs in that it allows for reuse of values computed by sub-trees, i.e. the resulting algorithm is a graph rather than a tree. The image processing algorithm that a given chromosome represents can be thought of as a directed acyclic graph where the non-terminal nodes are primitive image processing operations, and the terminal nodes are individual image planes extracted from the multi-spectral image used as input. Our representation also differs in that the total number of nodes is fixed (although not all of these may actually be used in the final graph), and crossover is carried out directly on the linear representation.

We have restricted our “gene pool” to a set of useful primitive image processing operators (“genes”). These include spectral, spatial, logical, and thresholding operators. The set of morphological operators is restricted to function-set processing morphological operators, i.e., gray-scale morphological operators having a flat structuring element. The sizes and shapes of the structuring elements used by these operators are also restricted to a pre-defined set of primitive shapes, which includes the square, circle, diamond, horizontal cross and diagonal cross, and horizontal, diagonal, and vertical lines. The shape and size of the structuring element are defined by operator parameters. Other local neighborhood/windowing operators such as mean, median, etc., specify their kernels/windows in a similar way. We define scratch planes as blocks of memory for storing intermediate calculations within a candidate image-processing algorithm. Once “scratch” planes have been generated, GENIE is allowed to explore weighted sums, differences and ratios of data and scratch planes.

A single gene consists of an operator, plus a variable number of input arguments specifying from where input is read, output arguments specifying where output is to be written, and any additional parameters that might be required to specify how the specific operator works (e.g., the diameter and shape of a structuring element used in a morphological filter). The operators used in GENIE take one or more distinct image planes as input, and generally produce a single image plane as output. Input can be taken from any data plane in the training data image cube. Output is written to one of a number of scratch planes, temporary workspaces where an image plane can be stored. Genes can also take input from scratch planes, but only if that scratch plane has been written to by another gene positioned earlier in the chromosome sequence. We use a notation for genes[1,4] that is may be illustrated by an example: the gene `[ADDP rD0 rS1 wS2]` applies pixel-by-pixel addition to two input planes, read from data plane 0 and from scratch plane 1, and writes its output to scratch plane 2. Any additional required operator parameters are listed after the output arguments. Our genetic algorithm works on strings of gene text blocks of this kind: cross-over preserves individual gene blocks, while mutation can alter the contents of gene blocks.

Note that although all chromosomes have the same fixed number of genes, the effective length of the resulting algorithm may be smaller than this. For instance, an operator may write to a scratch plane that is then overwritten by another gene before anything reads from it. GENIE performs an analysis of chromosome graphs when they are created and only carries out those processing steps that actually affect the final result. Therefore, the fixed length of the chromosome acts as a maximum effective length.

2.3 Supervised Classification and Fitness Evaluation

Each candidate image-processing algorithm generates a number of intermediate feature planes (or "signature" planes). Complete classification requires that the image-processing algorithm produce a binary-valued output plane for the scene. It is possible to treat, e.g., the contents of the first scratch plane as the final output for that candidate image- processing algorithm (thresholding would generally be required to obtain a binary result, though GENIE can choose to apply its own Boolean thresholding functions). However, we have found it to be useful to perform the combination of the data and scratch planes using a non-evolutionary method, and have implemented a supervised classifier backend. To do this, we first select a subset of the scratch planes and data planes to be "signature" planes. For the present experiments, this subset consists of just the scratch planes. We then use the provided training data and the contents of the signature planes to derive the Fisher Discriminant, which is the linear combination of the signature planes that maximizes the mean separation in spectral terms between those pixels marked up as "true "and those pixels marked up as "false", normalized by the total variance in the projection defined by the linear combination. The output of the discriminant-finding phase is a real-valued single-plane "answer" image. This is reduced to a binary image by exhaustive search over all the training pixels to find the threshold value that minimizes the total number of misclassifications (false positives plus false negatives) on the training data.

The fitness of a candidate solution is given by the degree of agreement between the final binary output plane and the training data. This degree of agreement is determined by the Hamming distance between the final binary output of the algorithm and the training data, with only pixels marked as true or false (as recorded in the weight plane) contributing towards the metric. The Hamming distance is then normalized so that a perfect score is 1000.

3. DETECTING CROPS WITH *HYPERION* HYPERSPECTRAL IMAGERY

GENIE was originally designed to evolve feature extraction algorithms for multispectral imagery, so the extention to hyperspectral imagery can be viewed as a test of the scalability of the technique and we increase the complexity of the data from 10's to 100's of spectral bands. Hyperspectral imagery is often analysed using purely spectral techniques such as the spectral angle mapper (SAM) algorithm, or by the design of matched filters that use the whole spectrum of each pixel (for a general review of hyperspectral image processing, see, e.g., Ref. 13). GENIE's set of primitive image processing operators from which it builds its candidate algorithms are designed to work on only one or two spectral bands of data each, and a complete candidate algorithm consisting of, e.g., 10 image processing steps might only use 10 to 15 bands of data from the hyperspectral datacube. Thus, the present experiment tests the ability of the GENIE system

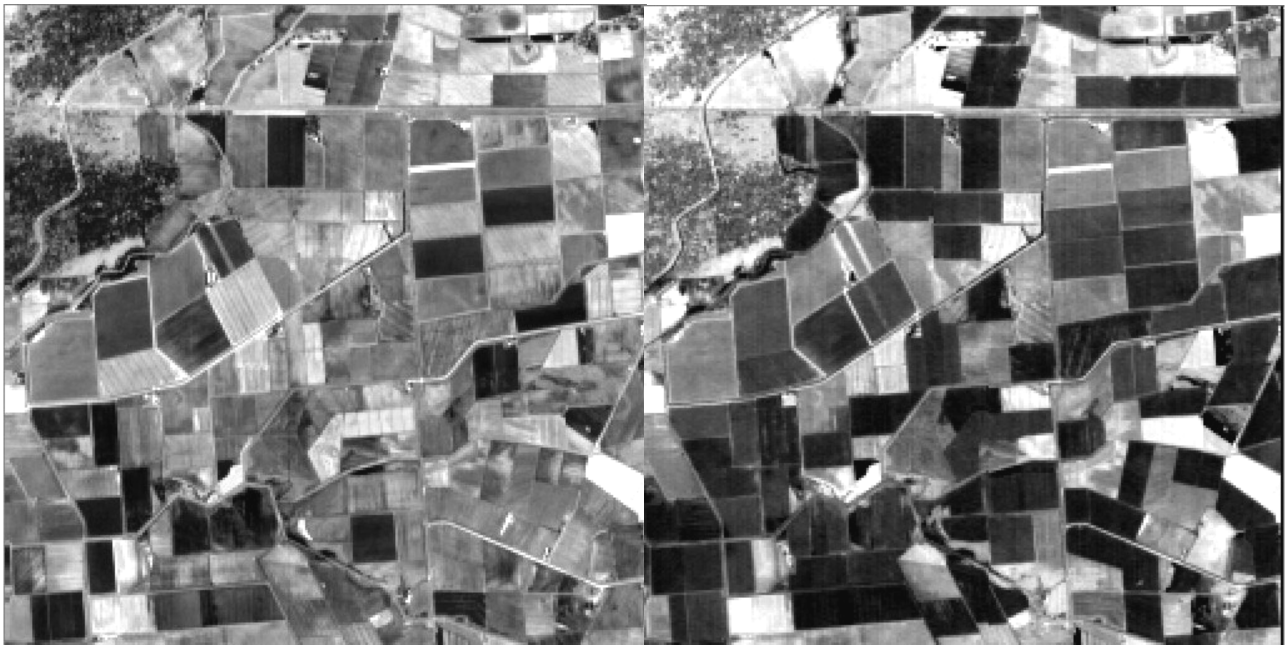

Figure 1. Visible (left) and Infrared (right) views of the training region extracted from the Hyperion sample scene. Rice, soy, and corn fields are present, as well as unplanted ploughed fields, natural vegetation, and roads and buildings. At 30m spatial resolution, textural differences between crops are noticeable.

to exploit the inherent redundancy of hyperspectral imagery and identify a small number of relevant bands out of the hyperspectral data cube.

Our hyperspectral imagery data source is a sample scene released by the Hyperion instrument team. Hyperion (see, e.g., Ref. 14 and references therein) is an experimental, moderate-resolution (~30m/pixel), 220 band visible (~0.4 μm) to short wave infrared (~2.5 μm) hyperspectral imager flown on the NASA New Millenium Program's Earth Observing 1 (EO-1) spacecraft. The scene we used covers part of the Coleambally Irrigation Area, an agricultural region located in the state of New South Wales, Australia (image collected 6 March, 2001). This region produces a number of commercial crops, including rice, corn and soy beans. Individual fields are large enough that even at ~30m/pixel spatial resolution, textural cues to the nature of planted crops are obvious (e.g., terracing of rice paddies). Figure 1 shows our 256x256 pixel training region, which is a small part of the full (417x752 pixel) sample scene. This region was chosen because of the availability of ground truth, in the form of a map of planted fields, shown in Fig. 2, that overlaps our training region.

4. RESULTS AND DISCUSSION

For each of three crops, rice, soy, and corn, a small amount of training data and a larger amount of testing data was marked up for the area shown in Figure 1, and the GENIE system was used to evolve a feature extraction algorithm for that crop (training data and result for each crop is shown in Fig. 3 ; example test data in shown in Fig. 4). Each run required approximately one hour of wall-clock time running on a standard Linux workstation. This is comparable to similar runs with multispectral imagery, and demonstrates no loss in computational performance with the more complex hyperspectral imagery. Table 1 presents our detection and false alarm rates for the image on the training data.

Feature	Training Scene		Test Scene	
	Detection Rate [%]	False Alarm Rate [%]	Detection Rate [%]	False Alarm Rate [%]
Rice	99.88	0.22	96.82	4.4
Soy	99.93	0.0	100.0	5.82
Corn	98.75	1.0	75.0	0.85

Table 1. GENIE result for crop detection on the training data

Considering qualitative performance outside of the training data, on comparing the results in Figure 3 with the ground truth in Figure 2 we see that the algorithms evolved by GENIE have marked out physically reasonable crop regions for each of the crops of interest. These agree well with the ground truth map of crop planting.

In the case of corn, fields in the south-west quadrant of field region #36 and north-west corner of field region #37 (see Fig.2) are not detected, but on inspection of the raw imagery (Fig.1) it is not clear that these fields have been correctly labeled in the "ground truth" (Fig.2). We intend to obtain clarification on this point from the Coleambally Irrigation Area managers. The result for soy outside the training area is good apart from the incorrect detection of field of sorghum (no examples of sorghum were used in the training phase). The result for rice is generally good.

The algorithms found by GENIE used both spatial and spectral processing. We will describe the soy algorithm in a little detail, as this is representative of our other results. Further details on all these algorithms, and comparisons to other standard algorithms, will appear in a future publication.

The soy algorithm appears in our genetic programming representation as the text string (see Refs. 1,4),

[SPIKE rD175 wS1 0.34 0.85][MSAVI rD17 rD115 wS2][R5R5 rD217 wS3] .

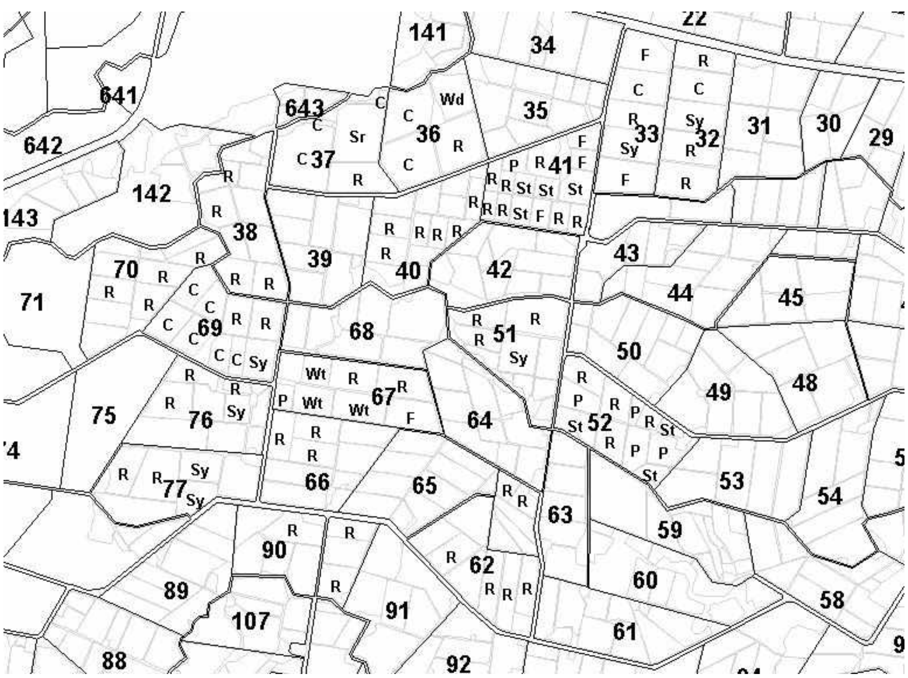

Figure 2. Ground truth for the fields shown in Fig. 1. Labels of interest are R: Rice, C: Corn, Sy: Soy. Proceedings of the Hyperion Data User Workshop, 2001.

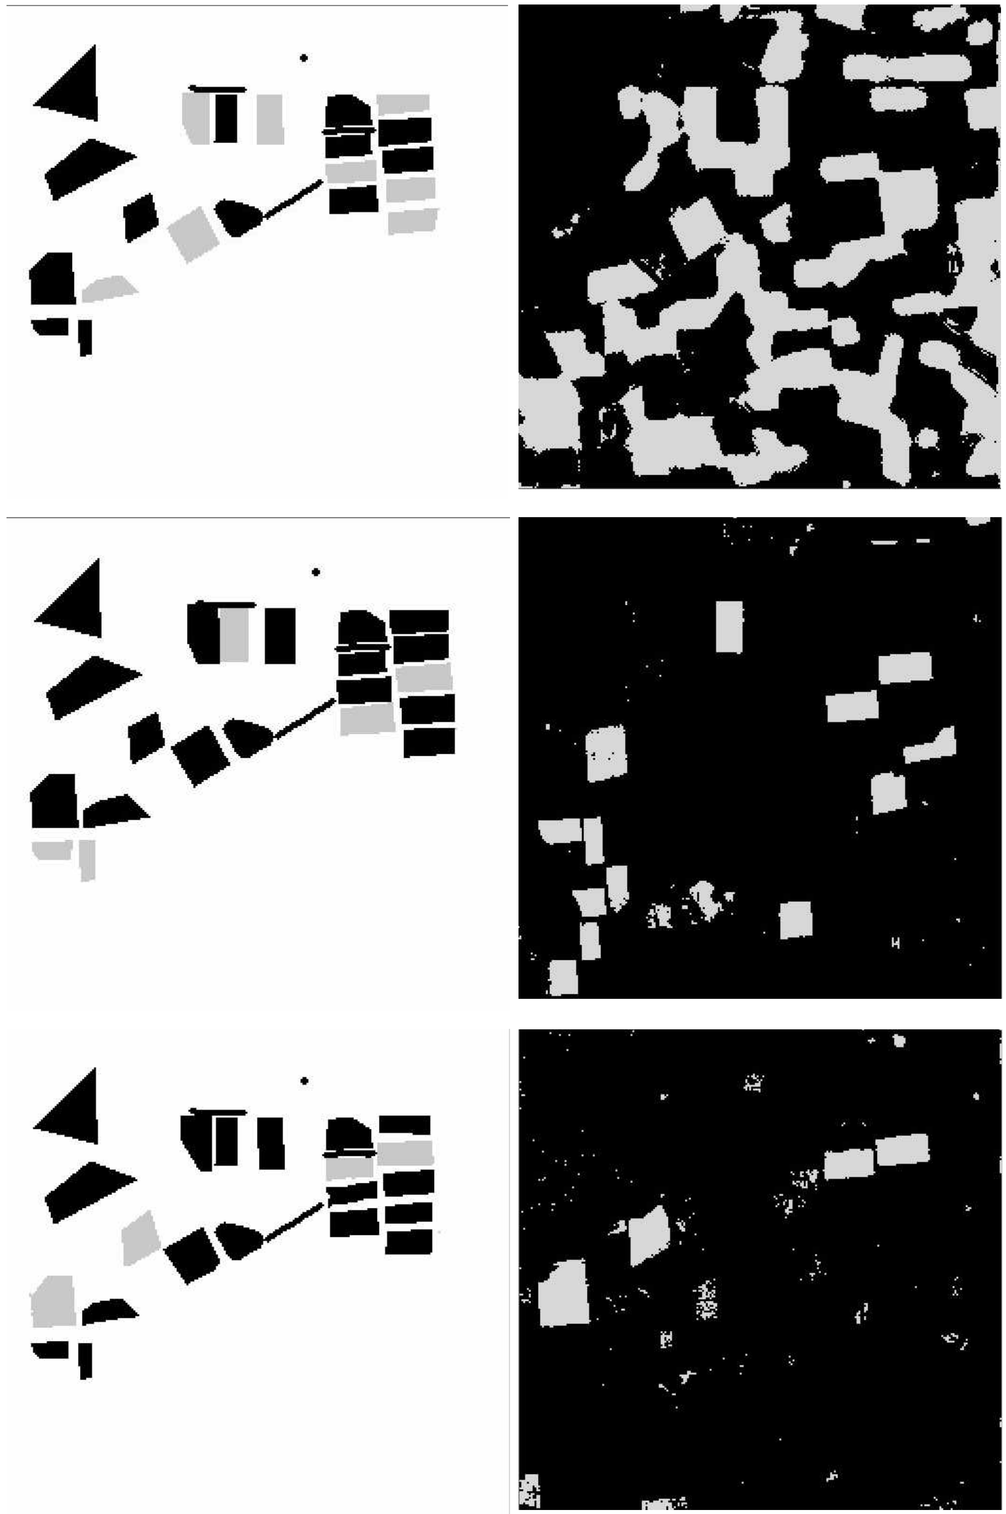

Figure 3. GENIE training data (left column) and results (right column) for rice, soy, and corn crops (from top to bottom). Grey pixels mark the feature of interest, and black pixels mark background. These results compare well to the ground truth presented in Figure 2.

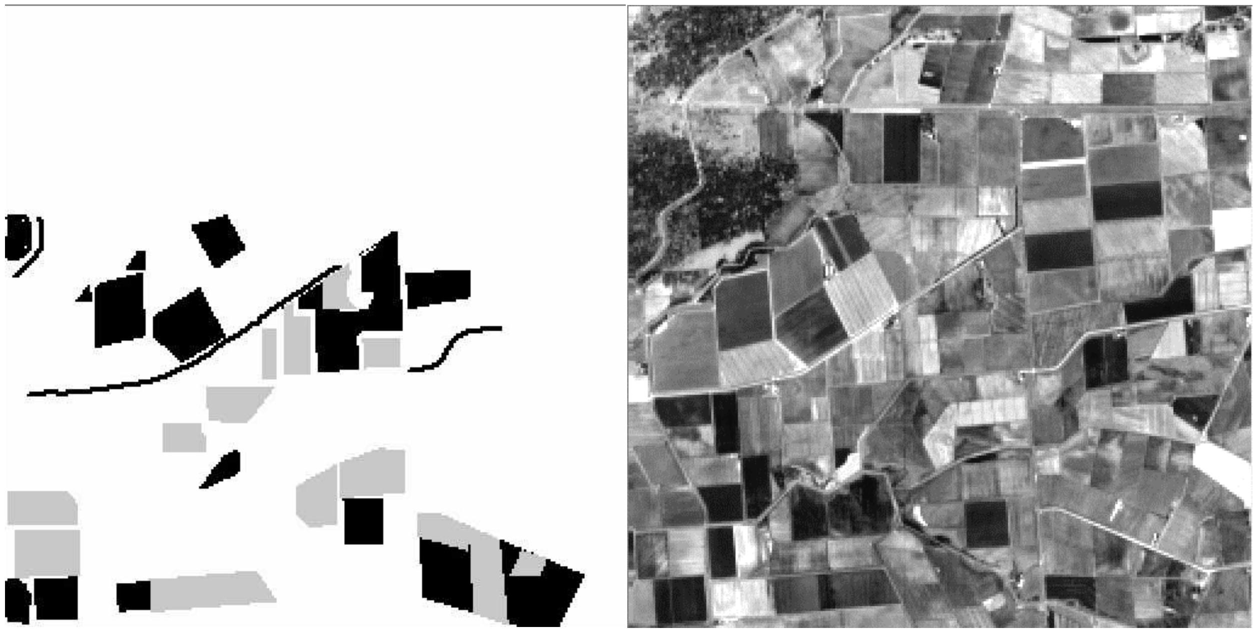

Figure 4. Testing data for rice fields. Grey pixels mark the feature of interest, and black pixels mark background.

This image processing algorithm constructs three spatio-spectral signature bands:

- An amplitude band-pass filter (SPIKE) is applied to spectral band 175 (1.901 μm), which passes pixels with values in a certain range, and sets pixels outside of that range to zero.

- A modified soil-adjusted vegetation index (MSAVI) function (see, e.g., Ref. 13) is applied using data bands 17 (0.519 μm) and 115 (1.296 μm).

- A local texture measure (which we call R5R5, see Ref. 4) is applied to data band 217 (2.325 μm).

A Fisher linear discriminant supervised classifier then finds the optimal linear combination of these bands, given the training data, and a boolean decision threshold is found to maximize detections and minimize false alarms.

From this, we can see that GENIE has in this case been able to identify a small (~1%) relevant subset of the hyperspectral datacube (visible to short-wave infrared), and has identified a useful mixture of spatial (R5R5) and spectral (SPIKE, MSAVI) processing on those planes. This result, and comparison of it to the results for rice, corn, and other crops, is now of interest for understanding the spatio-spectral signatures of these crops given this type of imagery, and for the identification of useful subsets of the hyperspectral range of wavelengths. For example, a multi-spectral imager could be tuned to exploit that features.

5. CONCLUSIONS

We have demonstrated evolution of spatio-spectral algorithms on hyperspectral imagery. In seeking to evolve algorithms to extract a range of agricultural crops, we find that the system was able to produce algorithms that distinguish between reasonably similar classes of vegetation, and that continue to perform well outside the training data. We find these results encouraging for future efforts of discovering hyperspectral signatures of vegetation.

ACKNOWLEDGEMENTS

We wish to acknowledge our colleagues on the ISIS/GENIE team from Los Alamos National Laboratory: Jeffrey J. Bloch, Nancy A. David, Diana Esch-Mosher, Mark Galassi, Neal R. Harvey, Reid B. Porter, Simon Perkins, John J. Szymanski, and James Theiler. We would like to thank the Hyperion science team for making available the hyperspectral imagery sample data used in this work. This work was supported by the Department of Energy and Department of Defense.

REFERENCES

1. S. P. Brumby, J. Theiler, S.J. Perkins, N. R. Harvey, J.J. Szymanski, J.J. Bloch, and M. Mitchell, "Investigation of feature extraction by a genetic algorithm", Proc. SPIE, Vol. 3812, pp. 24-31,1999.
2. J. Theiler, N. R. Harvey, S. P. Brumby, J. J. Szymanski, S. Alferink, S. J. Perkins, R. Porter, and J. J. Bloch, "Evolving retrieval algorithms with a genetic programming scheme", Proc. SPIE, Vol. 3753, pp.416-425, 1999.
3. S. Perkins, J. Theiler, S. P. Brumby, N. R. Harvey, R. B. Porter, J. J. Szymanski, and J. J. Bloch, "GENIE: A hybrid genetic algorithm for feature classification in multi-spectral images", Proc. SPIE, Vol. 4120, pp. 52-62, 2000.
4. N. R. Harvey, J. Theiler, S. P. Brumby, S. Perkins, J. J. Szymanski, J.J. Bloch, R.B. Porter, M. Galassi, and A. C. Young, "Comparison of GENIE and conventional supervised classifiers for multispectral image feature extraction", IEEE Transactions on Geoscience and Remote Sensing, Vol. 40, pp. 393-404, February 2002.
5. J. H. Holland, *Adaptation in Natural and Artificial Systems*, University of Michigan, Ann Arbor (1975).
6. I. Rechenberg, *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*, Fromman-Holzboog, Stuttgart (1973).
7. L. Fogel, A. Owens and M. Walsh, *Artificial Intelligence through Simulated Evolution*, Wiley, New York (1966).
8. J. R. Koza, *Genetic Programming: On the Programming of Computers by Natural Selection*, MIT, Cambridge (1992).
9. N. R. Harvey, S. Perkins, S. P. Brumby, J. Theiler, R. B. Porter, A. C. Young, A. K. Varghese, J .J. Szymanski, and J. J. Bloch, "Finding golf courses: The ultra high tech approach", Proc. Second European Workshop on Evolutionary Computation in Image Analysis and Signal Processing (EvoIASP2000), Edinburgh, UK, pp. 54-64, 2000.
10. S. P. Brumby, N. R. Harvey, J. J. Bloch, J. Theiler, S. Perkins, A. C. Young, and J. J. Szymanski, "Evolving forest fire burn severity classification algorithms for multi-spectral imagery", Proc. SPIE, Vol. 4381, pp. 236-245, 2001.
11. S. P. Brumby, J. Theiler, J. J. Bloch, N. R. Harvey, S. Perkins, J. J. Szymanski, and A. C. Young, "Evolving land cover classification algorithms for multi-spectral and multi-temporal imagery", Proc. SPIE, Vol. 4480, pp. 120-129, 2002.
12. Catherine S. Plesko, Steven P. Brumby and Conway Leovy, "Automatic Feature Extraction for Panchromatic Mars Global Surveyor Mars Orbiter Camera Imagery", Proc. SPIE, Vol. 4480, pp. 139-146, 2002.
13. J. A. Richards and X. Jia, *Remote Sensing Digital Image Analysis*, 3rd ed., Chapter 13, Springer, Berlin (1999).
14. Peter J. Jarecke, Karen E. Yokoyama, Pamela Barry, "On-orbit solar radiometric calibration of the Hyperion instrument", Proc. SPIE, Vol. 4480, pp. 225-230, 2002.
15. R. S. Lunetta and C. D. Elvidge (editors), *Remote sensing change detection*, Ann Arbor, Chelsea (1998).
16. The ENVI software package is described on the web site: http://www.rsinc.com/envi.